The Sovereign Model, The Silent Exclusion: How Government-Built Language Models Are Being Deployed in Jurisdictions With the Weakest International Disability Law Protections

An Investigative Exposé on the Collision Between State AI Sovereignty and the Global Disability Rights Architecture

---

TL;DR

A global race to build sovereign large language models (LLMs) is accelerating across the developing world, from Saudi Arabia's **ALLaM** (500 billion Arabic tokens) to the UAE's **Falcon**, Kazakhstan's **KazLLM**, and national AI strategies in **Nigeria, Egypt, Kenya, Indonesia, Vietnam, and Thailand**. Yet the jurisdictions making the largest public investments in government-built AI are, in nearly every case, the same jurisdictions with the **weakest international disability law protections** — measured by CRPD implementation gaps, the non-ratification of the African Disability Protocol by 75% of AU member states, absent ICT accessibility standards, and no mandatory accessibility criteria in public procurement. The result is a systematic structural failure: **state AI systems are being deployed as public infrastructure without any enforceable obligation to serve the 1.3 billion people worldwide with disabilities**, and without any disability-disaggregated evaluation of model performance. This exposé documents the scale of that failure, the legal vacuum enabling it, and the specific governments where the risk of algorithmic exclusion is most acute.

---

1. The Structural Contradiction: AI Sovereignty Without Accessibility Sovereignty

1.1 The Global Surge in Government-Built Language Models

The period from 2023 to 2026 has witnessed an unprecedented expansion of state-directed artificial intelligence development. Governments across every inhabited continent have launched national LLM initiatives, driven by a convergence of geopolitical competition, digital sovereignty concerns, and the strategic imperative to reduce dependence on foreign — primarily American and Chinese — AI infrastructure. Saudi Arabia's **SDAIA** mobilized **16 public entities** to aggregate a **500 billion-token Arabic dataset** for its **ALLaM** model, which now runs on both IBM Watsonx and Microsoft Azure in 7B, 13B, and 70B parameter configurations (IBM Newsroom - Middle East & Africa) . The UAE's **Technology Innovation Institute (TII)** released **Falcon 2 11B**, a multilingual model tested against Meta's Llama 3 and Google's Gemma, with a vision-to-language variant marketed as having "new horizons for accessibility in AI") (. Kazakhstan's **ISSAI** developed **KazLLM** to address the "AI gap for low-resource languages" through transfer learning and data augmentation techniques that could, in principle, be replicated by "Ethiopia (Amharic), Nigeria (Hausa, Yoruba), Thailand (Thai), and Indonesia (Javanese, Sundanese)" (Medium) .

Across Sub-Saharan Africa, the momentum is equally pronounced. Egypt launched its **Second National AI Strategy (2025–2030)** built around six pillars — governance, technology, data, infrastructure, ecosystem, and talent — with targets including **7.7% of GDP from ICT by 2030** and training **30,000 AI specialists** (ciphersense.ai) . Kenya adopted its **National AI Strategy 2025–2030** in

March 2025, focused on digital infrastructure, data ecosystems, and research & innovation, backed by an **AI 10 billion program** in partnership with UNDP and the African Development Bank (ciphersense.ai) . Nigeria revised its **National Artificial Intelligence Strategy (NAIS)** in September 2025 and launched an **AI Scaling Hub**, while Rwanda's **National AI Policy** emphasizes ethical AI with four strategic pillars (ciphersense.ai) . The African Union's **Continental AI Strategy**, adopted in June 2024, calls for "African AI" tailored to local needs, realities, and cultures (cipit.org) . By February 2025, **40 African countries** had enacted or strengthened data protection laws addressing AI-related privacy concerns (cipit.org) .

In Southeast Asia, the pattern replicates. Thailand's **National AI Strategy & Action Plan (2022–2027)** aligns with the Thailand 4.0 vision, deploying AI in agriculture, healthcare, and smart cities (ginc.org) . Vietnam's **National Strategy on R&D and Application of AI to 2030** sets targets to rank among the **top four ASEAN AI leaders**, training **20,000 specialists** and fostering **50+ domestic AI enterprises** (ginc.org) . Indonesia's **National AI Strategy 2020–2045** embeds AI within the **Making Indonesia 4.0** industrial modernization agenda (ginc.org) . The Philippines launched its **National AI Strategy Roadmap 2.0 (2024)** to become a regional AI Center of Excellence (ginc.org) . Even Myanmar, under military rule, has convened **four coordination meetings** to draft a National AI Strategy and Policy (昆明信息港) . Turkmenistan, one of the world's most politically isolated states, is now developing a National AI Strategy with UNDP support, built on its **State Programme for the Digital Economy 2026–2028** (United Nations Development Programme) .

1.2 The Accessibility Governance Gap

The defining structural feature of this global sovereign AI movement is not its scale or ambition — it is the **near-total absence of disability accessibility as a design requirement, procurement criterion, or evaluation metric** in every single national AI strategy and government LLM program documented above. Not one of the national AI strategies referenced — Egypt's six-pillar framework, Kenya's three-pillar plan, Nigeria's NAIS, Thailand's 2022–2027 action plan, Vietnam's 2030 strategy, Indonesia's 2020–2045 roadmap, or the Philippines' Roadmap 2.0 — contains a dedicated disability accessibility provision, a requirement for inclusive dataset construction, or a mandate for disability-disaggregated model evaluation (cipit.org) . The AU's Continental AI Strategy, despite its emphasis on "African AI" tailored to local needs, makes no specific reference to the **1.3 billion people with disabilities worldwide** or the estimated **80 million people with disabilities in Africa alone** (cipit.org) .

This is not an oversight. It is the predictable outcome of a structural governance gap in which **AI policymaking and disability rights policymaking operate in entirely separate institutional silos**, even within the same governments. Nigeria's National Commission for Persons with Disabilities (NCPWD) and its National Centre for AI and Robotics (NCAIR) report through different ministerial channels, with no formal coordination mechanism linking disability rights obligations to AI procurement decisions (NCPWD) . Saudi Arabia's SDAIA, which built ALLaM, operates under the National Information Center, while the Authority for the Care of Persons with Disabilities functions as a separate entity under a different ministry — and the CRPD Committee has already noted that Saudi accessibility initiatives are "**not systematic**, are **often associated with charitable initiatives**, and are **limited in scope**" (International Disability Alliance) . Egypt, which ratified the CRPD in 2008 and launched its Second National AI Strategy in January 2025, has **not ratified the African Disability Protocol** and has no comprehensive national accessibility legislation covering ICT (Centre for Human Rights) .

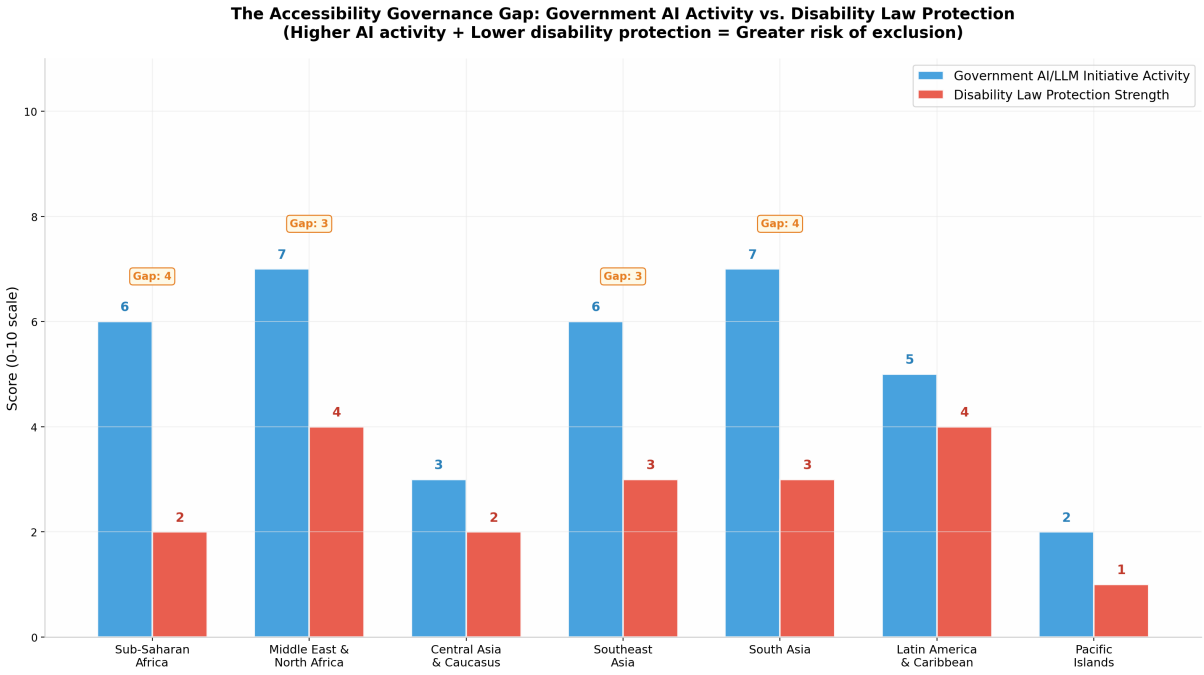

Figure 1: Governance Gap Chart

Figure 1: The Accessibility Governance Gap across regions. The divergence between government AI activity (blue) and disability law protection strength (red) represents the structural risk of algorithmic exclusion. Sub-Saharan Africa and South Asia show the widest gaps, where ambitious AI strategies coexist with minimal disability rights frameworks.

---

2. The Legal Architecture of Exclusion: Where International Disability Law Is Weakest

2.1 The CRPD's Promise and the Implementation Chasm

The **Convention on the Rights of Persons with Disabilities (CRPD)**, which entered into force in May 2008, represents the most comprehensive international human rights instrument addressing disability. **Article 9** explicitly requires States Parties to "take appropriate measures to ensure to persons with disabilities access, on an equal basis with others, to… information and communications, including **information and communications technologies and systems**" (un.org) . Article 9(2)(g) mandates states to "promote access for persons with disabilities to **new information and communications technologies and systems**, including the Internet," while Article 9(2)(h) requires promoting "the design, development, production and distribution of **accessible information and communications technologies and systems at an early stage**, so that these technologies and systems become accessible at minimum cost" (un.org) . These provisions are not optional guidance — they are binding treaty obligations on the **185 states** that have ratified the Convention (OHCHR) .

Government-built LLMs squarely fall within the scope of Article 9. A state-developed language model deployed for public services — whether as a chatbot for government portals, an AI-powered document processing system, a healthcare diagnostic assistant, or an educational tutoring tool — is an "information and communications technology and system" provided by the state. Under CRPD Article 4(1)(b), states must "adopt all appropriate legislative, administrative and other measures for the implementation of the

rights recognized in the present Convention" (OHCHR) . Article 4(1)(f) requires states to "undertake or promote research and development of universally designed goods, services, equipment and facilities" (era-comm.eu) . The CRPD Committee's **General Comment No. 2 (2014)** on accessibility clarifies that these obligations extend to "all ICT products and ICT based applications and services, a far-reaching implication for industry, governments and civil society" (G3ict) .

Yet the CRPD Committee's **concluding observations** on state party reports reveal a systematic pattern of non-compliance with Article 9 in precisely those jurisdictions now launching AI initiatives. The Committee expressed concern to the **United Arab Emirates** that "almost all initiatives to ensure accessibility for persons with disabilities address physical disabilities, are **not systematic**, are **often associated with charitable initiatives** and are **limited in scope**," and noted "the **absence of specific binding legal and policy frameworks** to ensure accessibility… including access to information, means of communications and transport" (International Disability Alliance) . To **Indonesia**, the Committee noted "the **lack of comprehensive national accessibility legislation and strategy** across all levels of government, includes accessibility standards, monitoring and enforcement mechanisms and that covers all areas, including… **information and communications technology**" (International Disability Alliance) . To the **Philippines**, it observed that existing laws "**do not include principles of accessibility through universal design**" and that there is "**no national action plan for accessibility**" (International Disability Alliance) . To **Thailand**, the Committee raised concerns about the lack of "comprehensive and effective measures… including the **lack of information and communication technologies and systems**" (International Disability Alliance) .

2.2 The African Disability Protocol: A Continental Failure

The **Protocol to the African Charter on Human and Peoples' Rights on the Rights of Persons with Disabilities in Africa** (African Disability Protocol, or ADP) was adopted by the African Union in 2018 to provide a continent-specific framework for disability rights protection. It addresses African-specific contexts — including the relationship between disability and poverty, the impact of conflict on disability, and the role of traditional practices — that the CRPD does not fully capture. The ADP requires **15 ratifications** to enter into force. As of 2025, **only 14 states** have ratified and deposited their instruments: Angola, Burundi, Cameroon, Kenya, Mali, Malawi, Mozambique, Namibia, Niger, Nigeria, Rwanda, Sahrawi Arab Democratic Republic, South Africa, and Uganda (Centre for Human Rights) .

Forty-one of 55 AU member states — 75% of the continent — have not ratified the ADP. Among the non-ratifiers are some of Africa's most AI-ambitious nations. **Egypt**, which leads the continent in AI readiness (51st globally) and has a comprehensive Second National AI Strategy, has **not signed the ADP** (ciphersense.ai) . **Ghana**, with an active AI ecosystem and a national AI strategy in development, has **not signed** (Centre for Human Rights) . **Ethiopia**, which finalized its national AI strategy in 2025, has **not signed** (Diplo) . **Mauritius**, ranked 71st globally in AI readiness with its Blueprint 2025–2029, has **not signed** (ciphersense.ai) . **Tunisia**, an early mover in AI governance with a National AI Governance Framework, has **not signed** (OECD) . **Senegal**, which is building a national data centre and supercomputer, has **not signed** (OECD) .

Even more critically, **25 AU member states have not even signed the ADP**, meaning they have not begun the domestic consultation and legislative processes required for ratification. These include countries with significant disability populations and minimal legal protections: **Somalia** (no disability-specific legislation), **South Sudan** (no disability legislation), **Eritrea** (no disability legislation until a CRPD ratification in January 2025), **Equatorial Guinea** (no disability legislation), **Guinea-Bissau**

(no disability legislation), **Liberia** (no disability legislation), **Madagascar** (no disability legislation), and **Mauritania** (no disability legislation) (Centre for Human Rights) .

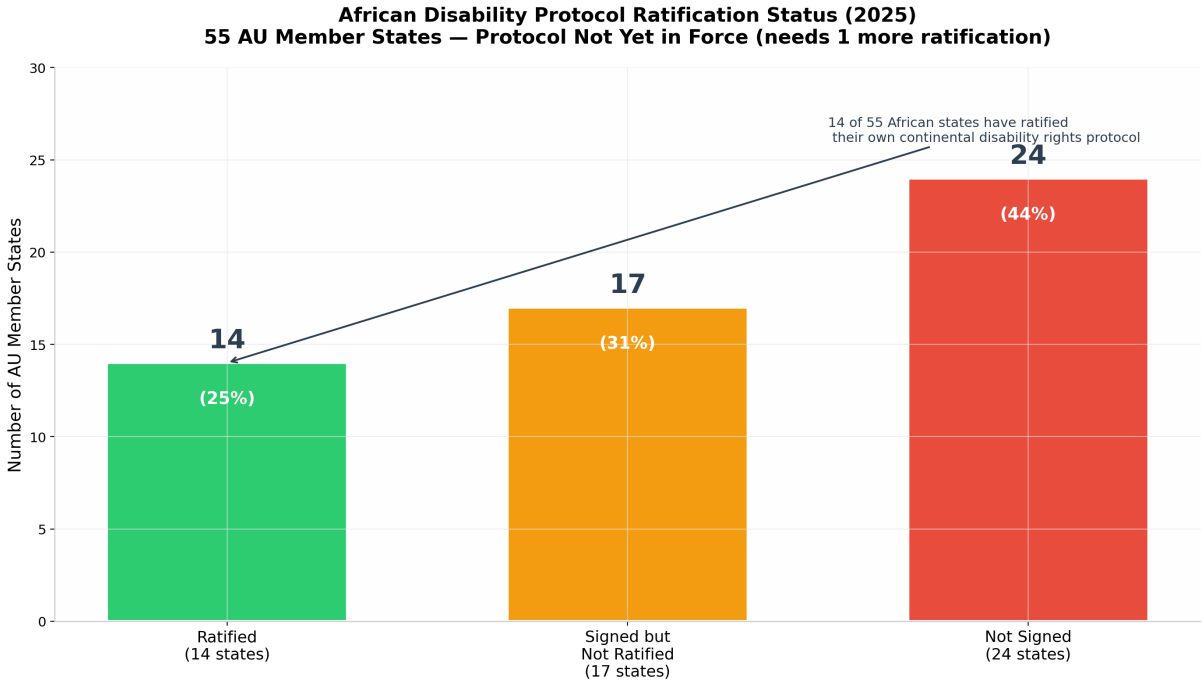

Figure 2: ADP Ratification Chart

Figure 2: The African Disability Protocol ratification status as of 2025. With only 14 of 55 AU member states having ratified, the Protocol remains one ratification short of entry into force. Twenty-four states — 44% of the continent — have not even signed the treaty.

2.3 The Marrakesh Treaty Gap and Accessible AI Formats

The **Marrakesh Treaty to Facilitate Access to Published Works for Persons Who Are Blind, Visually Impaired or Otherwise Print Disabled** (2013) creates a mandatory exception to copyright law allowing authorized entities to reproduce works in accessible formats without seeking permission from rights holders. As of 2026, **93 countries** have ratified or acceded to the Treaty (Association of Research Libraries) . The Treaty defines "accessible format" broadly to include "Braille, large print, adapted e-book, audio book, and any alternative way or form that gives beneficiaries access to the work" (Association of Research Libraries) .

The intersection between the Marrakesh Treaty and government AI systems has been **virtually unexplored in policy or legal scholarship**. When a state develops an LLM and deploys it as a public service — for example, as a document summarization tool for government publications, an educational content generator, or a legal information assistant — the outputs of that system must be accessible to persons with print disabilities. If the LLM generates content in formats that are not compatible with screen readers, does not produce alternative text for visual elements, or cannot output content in Braille-ready formats, the state is effectively reproducing the very barriers the Marrakesh Treaty was designed to eliminate. Yet no national AI strategy reviewed for this exposé references the Marrakesh Treaty, and no government LLM model card includes an accessibility format compliance statement.

---

3. Case Studies in Algorithmic Exclusion: Government LLMs in Low-Protection Jurisdictions

3.1 Saudi Arabia's ALLaM: 500 Billion Tokens, Zero Disability Data Points

Saudi Arabia's **ALLaM** (Arabic Large Language Model) represents one of the most technically ambitious and well-resourced government LLM projects in the developing world. Developed by the **National Center for Artificial Intelligence (NCAI)** at SDAIA, ALLaM was trained on a **500 billion-token Arabic dataset** aggregated from **16 public entities**, including the text of **300 Arabic language books**, and evaluated through **more than 1 million prompts** generated by **400 subject matter experts** (middleeastainews.com) . The model outperforms all other LLMs in Arabic-language queries on the MMLU (Massive Multitask Language Understanding) benchmark (middleeastainews.com) . ALLaM is available in 7B, 13B, and 70B parameter versions on both IBM Watsonx and Microsoft Azure (IBM Newsroom - Middle East & Africa) .

What the SDAIA documentation does **not** mention is any aspect of disability accessibility. The 500 billion-token training dataset, drawn from government publications, books, and public entity documents, was not filtered or annotated for accessibility content. The 400 subject matter experts who generated the 1 million evaluation prompts were not instructed to test the model's performance on disability-related queries. The model card for ALLaM does not include any disability-disaggregated performance metrics. When SDAIA's Director Dr. Esam Alwagait stated that the collaboration with IBM would "serve as a catalyst for further technological advancements" and "propel the region into a new era of digital excellence," he did not mention the **1.3 million people with disabilities in Saudi Arabia** who may be unable to access services built on this infrastructure (IBM Newsroom - Middle East & Africa) .

A 2025 qualitative study of AI healthcare experiences among people with disabilities in Saudi Arabia found that "**current platforms often lack screen reader compatibility, voice navigation options, and disability-specific user interface adaptations**" and that "**the Sehhaty application... presents navigation challenges for users with visual impairments**" (MDPI) . The Sehhaty platform is part of the same Vision 2030 digital ecosystem that produced ALLaM. A participant with a speech impairment reported that "**the AI speech recognition system struggles with my speech patterns. Despite being marketed specifically for speech impairments, it requires constant retraining to understand my voice**" (MDPI) . Another participant with a visual impairment noted that "**the interface isn't compatible with my screen reader, so I need assistance to interpret the directions**" (MDPI) . These findings from Saudi Arabia's existing AI healthcare deployments raise serious questions about whether ALLaM, trained on the same institutional data and evaluated by the same expert panels, would perform any differently.

The CRPD Committee's concluding observations on Saudi Arabia noted that "**almost all initiatives to ensure accessibility for persons with disabilities address physical disabilities**" and that there is an "**absence of specific binding legal and policy frameworks**" for ICT accessibility (International Disability Alliance) . Saudi Arabia's GDPPI score of **88/100** — the second-highest among developing nations — reflects strong AI strategy inclusion through Vision 2030 and the establishment of the Authority for the Care of Persons with Disabilities, but the "representation" component remains "emerging" as new institutions develop capacity (policyjssr.com) .

3.2 The UAE's Falcon: Marketing Accessibility Without Measuring It

The UAE's **Falcon 2 11B VLM** (Vision-to-Language Model) represents a different but equally instructive case. TII's launch announcement explicitly highlighted the model's "**vision-to-language capabili-**

ties" as opening "**new horizons for accessibility in AI, empowering users with transformative image to text interactions**") (. The model can "**identify and interpret images and visuals from the environment**," with applications across "healthcare, finance, e-commerce, education and legal sectors" ranging from "**document management, digital archiving and context indexing to supporting individuals with visual impairments**") (.

This framing illustrates a pattern common across government AI initiatives in low-protection jurisdictions: **accessibility is invoked as a marketing claim, not enforced as a design requirement**. The Falcon 2 model card includes performance benchmarks against Meta's Llama 3 and Google's Gemma on standard NLP tasks, but it does not include any evaluation of the model's performance on accessibility-specific tasks — such as the accuracy of image-to-text conversion for assistive technology users, the model's handling of disability-related terminology in Arabic and English, or its ability to generate alt-text descriptions that meet WCAG 2.1 standards) (. The claim that the model supports "individuals with visual impairments" was not validated through testing with blind or low-vision users. No disability organization was involved in the model's development or evaluation.

The UAE ratified the CRPD in 2010 and has Federal Law No. 29 of 2006 on the Protection of People with Disabilities, which requires accessible digital services (Deque Systems) . Dubai's Digital Accessibility Policy (2020) enforces **WCAG 2.1 AA** for government entities (Deque Systems) . However, the CRPD Committee found that UAE accessibility initiatives remain limited in scope and "**associated with charitable initiatives**" rather than being embedded as enforceable rights (International Disability Alliance) . The gap between Dubai's WCAG 2.1 AA mandate for government websites and the absence of any accessibility evaluation for Falcon 2 — a government-funded model that will power those same websites — exemplifies the institutional disconnect between digital accessibility policy and AI development practice.

3.3 Nigeria's AI Strategy: 45/100 on Disability Inclusion

Nigeria presents perhaps the most stark illustration of the governance gap. Africa's most populous nation and largest economy, Nigeria revised its **National Artificial Intelligence Strategy (NAIS)** in September 2025 and launched an **AI Scaling Hub** through the National Centre for AI and Robotics (NCAIR) (ciphersense.ai) . The country's AI ecosystem includes **120+ AI startups** and **$1.2 billion in cumulative funding**, with the largest developer community on the continent (ciphersense.ai) . Yet Nigeria's **GDPPI score of 45/100** places it in the "high risk" zone for disability exclusion from AI systems (policyjssr.com) .

The reasons for this low score are structural. Nigeria ratified the CRPD in 2007 and passed the **Discrimination Against Persons with Disabilities (Prohibition) Act 2018**, which recognizes the right to equal employment opportunities and recommends a **5% quota** for persons with disabilities in public organizations (NCPWD) . But **only 23 of 36 states** have enacted corresponding state-level legislation (globaldisabilityfund.org) . The National Commission for Persons with Disabilities (NCPWD), established in 2020, is "**perceived as confused, often competing with disability-focused CSOs for implementation roles**" and its "engagement with OPDs [Organizations of Persons with Disabilities] is weak" (globaldisabilityfund.org) . The NCPWD's OPD Support Grant of NGN 20 million ($12,500) in 2023 "**has yet to be disbursed**" as of 2024 (globaldisabilityfund.org) .

Nigeria ratified the African Disability Protocol in October 2023 — a significant step — but the Protocol is not yet in force continent-wide, and Nigeria's domestic disability infrastructure remains critically underresourced. A 2024 study on financial inclusion for persons with disabilities in Nigeria found that "**persons with disabilities are confronted with multidimensional difficulties due to the inadequacy of laws to advance financial access to them**" and that existing guidelines require only "**two per cent**

of ATMs" to be equipped with tactile symbols for visually impaired clients (up.ac.za) . The country's AI strategy does not mention disability accessibility, inclusive dataset construction, or the participation of persons with disabilities in AI governance.

3.4 Kazakhstan's KazLLM: Evaluated on Legal Accuracy, Not Accessibility

Kazakhstan's **KazLLM**, developed by the **Institute of Smart Systems and Artificial Intelligence (ISSAI)** at Nazarbayev University, offers a revealing case of how evaluation frameworks for government LLMs systematically exclude disability metrics. A 2025 study evaluated KazLLM alongside GPT-4o mini, GEMMA, LLaMA, Phi, Qwen, and Mistral on a legal question-answering task for Kazakh legislation, using **ROUGE, METEOR, and F1 metrics** plus expert evaluation on five criteria: **legal accuracy, completeness, clarity, relevance, and professional style** (MDPI) . The study found that GPT-4o mini outperformed all models, while KazLLM "**faced serious difficulties in processing complex legal syntax and semantics**" and exhibited "**a tendency to hallucinations**" (MDPI) .

Not one of the seven evaluation criteria — nor any of the three quantitative metrics — addressed accessibility. The study did not test whether KazLLM's outputs could be parsed by screen readers, whether the model could generate responses in formats compatible with assistive technologies, or whether it performed differently when queried about disability rights under Kazakh law. The planned "online platform accessible to a wide range of users" — legal consultants, civil servants, and citizens — does not include accessibility requirements in its design specification (MDPI) . Kazakhstan ratified the CRPD in 2015 and ratified the Marrakesh Treaty in 2020, but its disability rights framework does not extend to ICT accessibility standards for government AI systems.

---

4. The Data Desert: Why Government LLMs Cannot Serve Disabled Users

4.1 The Absence of Disability Data in Training Corpora

The fundamental technical reason government LLMs fail disabled users is not malice or negligence — it is **data absence**. Large language models learn from the statistical patterns in their training data. If that data does not contain sufficient examples of disability-related content, the models will not generate accurate, relevant, or appropriate responses when queried about disability topics. The 2025 **AccessEval** benchmark study, which systematically evaluated disability bias across major LLMs, found "**systematic bias in disability-aware queries, with higher negative social perception, increased negative sentiment, and more factually incorrect and denial responses**" (ACL Anthology) . People with disabilities "**often receive less useful [responses] compared to normal users**," with the greatest performance drops in **finance, hospitality, and education** domains and the most affected disability types being **mobility, speech, and hearing impairments** (ACL Anthology) .

The training datasets for government LLMs compound this problem. SDAIA's ALLaM was trained on "**the text of 300 Arabic language books**" and aggregated government documents from 16 public entities (middleeastainews.com) . Falcon was trained on "**5.5 trillion tokens**" of multilingual web data) (. KazLLM was trained on legislative corpora from Kazakh legal portals (MDPI) . None of these datasets were constructed with disability content inclusion as a criterion. The Masakhane initiative, which has built the most comprehensive open datasets for African languages, does not include disability-specific annotations (cipit.org) . Digital Umuganda's Open Data for All (OD4A) initiative in Rwanda focuses on "**voice recordings, text samples, and translations**" for African languages but does not include disability-specific data collection (cipit.org) . Kencorpus in Kenya, GhanaNLP in Ghana, and

NajjaVoices in Nigeria advance NLP for underrepresented African languages, but none explicitly address disability representation (cipit.org) .

A 2023 study from Penn State's Human Language Technologies Lab found that "**all of the public models we studied exhibited significant bias against disability**" and that there was "**a problematic tendency to classify sentences as negative and toxic based solely on the presence of disability-related terms, such as 'blind,' without regard for contextual meaning**" (Pennsylvania State University) . The researchers created the **BITS (Bias Identification Test in Sentiment)** corpus to help identify explicit disability bias in AI sentiment analysis models, finding that "**statements referring to people with disabilities versus other control categories received significantly more negative and toxic scores**" (Pennsylvania State University) .

4.2 The Evaluation Blind Spot

The second technical failure is the **absence of disability-disaggregated evaluation** in every government LLM model card reviewed for this exposé. Model cards — the standardized documentation that accompanies published AI models — typically include accuracy on standardized benchmarks, sometimes broken down by language or domain, "**very rarely broken down by user characteristics, and almost never broken down by disability**" (Human in the Loop) . The ImageNet benchmark, which became the standard for computer vision research, "**did not report performance by subject characteristics**" (Human in the Loop) . The GLUE benchmark for natural language understanding "**did not either**" (Human in the Loop) .

The 2025 AccessEval framework represents the first systematic attempt to benchmark disability bias in LLMs, introducing metrics for "**response quality degradation**" and "**social perception biases**" across disability types (ACL Anthology) . But AccessEval has not been adopted by any government LLM program. Saudi Arabia's ALLaM was evaluated on MMLU — a general-knowledge benchmark that does not include disability-related questions (middleeastainews.com) . The UAE's Falcon was evaluated against Llama 3 and Gemma on standard NLP tasks) (. Kazakhstan's KazLLM was evaluated on ROUGE and METEOR for legal question-answering (MDPI) .

A 2025 paper by Waters of the Cognitive and Neurodiversity AI Lab at Virginia State University argued that "**there is no standardised methodology for evaluating whether AI systems work for disabled users, and the dominant evaluation practices actively disguise their failure to do so**" (Human in the Loop) . Accuracy and generalisability, the two metrics that have governed machine-learning benchmarks for a decade, "**treat disabled users as statistical noise rather than as a population whose performance matters**" (Human in the Loop) .

---

5. The Procurement Void: How Governments Buy AI Without Accessibility Requirements

5.1 The Global Procurement Gap

Public procurement — the process by which governments purchase ICT products and services — is the single most powerful lever for ensuring accessibility at scale. When a government mandates accessibility criteria in its procurement specifications, vendors must comply to win contracts, creating market-wide incentives for inclusive design. When accessibility is omitted from procurement, vendors have no commercial reason to invest in it, and accessibility becomes an afterthought at best.

The research on this point is unequivocal. A comprehensive study of ICT accessibility criteria in public procurement across OECD countries found that "**voluntary government schemes were not successful**" and that "**mandatory processes based on uniform global standards coupled with compliance will have an impact**" (Academia.edu) . **Only the USA and Japan** have comprehensive mandatory accessibility criteria in public procurement legislation (Academia.edu) . South Korea achieves over 90% compliance with web accessibility standards through a voluntary approach, but this is an outlier driven by unique cultural and institutional factors (Academia.edu) .

For the developing countries launching government AI initiatives, the procurement situation is substantially worse. The WEF/ITU Model ICT Accessibility Policy Report, adapted for city-wide procurement, recommends that governments "**clearly document and ensure the city's compliance with national legislation on accessibility and procurement**" and "**define a clear approach for ensuring that the city develops inclusive and accessible smart city tech**" even where national legislation does not exist (World Economic Forum) . But no national AI strategy reviewed for this exposé includes accessibility procurement requirements. Egypt's Second National AI Strategy, with its six pillars and 30,000-AI-specialist training target, does not mention procurement (ciphersense.ai) . Kenya's National AI Strategy 2025–2030, focused on digital infrastructure and data ecosystems, does not mention procurement (ciphersense.ai) . Nigeria's NAIS, with its AI Scaling Hub and 120+ startups, does not mention procurement (ciphersense.ai) .

5.2 The CRPD Committee's Procurement Directives

The CRPD Committee has repeatedly directed states to incorporate accessibility into public procurement, and its concluding observations provide a roadmap for what compliant procurement would look like. To the **UAE**, the Committee recommended "**obligatory accessibility criteria for public procurement**" as part of a comprehensive accessibility plan (International Disability Alliance) . To **North Macedonia**, it recommended "**accessibility standards as a requirement in all public procurement of services and construction**" (International Disability Alliance) . To **Poland**, it recommended "**accessibility through universal design in the Public Procurement Law**" (International Disability Alliance) . To **Hungary**, it recommended reviewing "**regulations on public procurement, and establish mandatory requirements and timelines for ensuring accessibility**" (International Disability Alliance) .

These recommendations have not been implemented in the jurisdictions now deploying government LLMs. Saudi Arabia's procurement of ALLaM development services from IBM and Microsoft did not include accessibility evaluation criteria (IBM Newsroom - Middle East & Africa) . The UAE's procurement of Falcon development through TII did not include disability-disaggregated benchmarking) (. Kazakhstan's development of KazLLM at Nazarbayev University did not include accessibility testing with assistive technology users (MDPI) . Nigeria's NCAIR procurement for its AI Scaling Hub does not include accessibility requirements for models or platforms (ciphersense.ai) .

---

6. Regional Analysis: The Geography of Algorithmic Exclusion

6.1 Sub-Saharan Africa: The Largest Gap

Sub-Saharan Africa presents the most acute case of the sovereign AI–disability rights collision. The region has the world's highest disability prevalence rates, the weakest legal protection frameworks, and the fastest-growing AI policy momentum — a combination that creates maximum exposure to algorithmic

exclusion.

Country	AI Strategy Status	ADP Ratification	CRPD Ratification	Disability Legislation	GDPPI Score
Nigeria	NAIS revised 2025, AI Scaling Hub (ciphersense.ai)	Ratified Oct 2023 (Centre for Human Rights)	Ratified Sep 2010 (OHCHR)	Prohibition Act 2018 (globaldisability-fund.org)	**45/100** (policyjssr.com)
Egypt	2nd National AI Strategy 2025–2030 (ciphersense.ai)	**Not signed** (Centre for Human Rights)	Ratified Apr 2008 (OHCHR)	Law 10/2018 (Centre for Human Rights)	**42/100** (policyjssr.com)
Kenya	National AI Strategy 2025–2030 (ciphersense.ai)	Ratified Nov 2021 (Centre for Human Rights)	Ratified May 2008 (OHCHR)	PWD Act 2003 (Centre for Human Rights)	**38/100** (policyjssr.com)
Ghana	AI strategy in development (Diplo)	**Not signed** (Centre for Human Rights)	Ratified Jul 2012 (OHCHR)	PWD Act 2006 (Centre for Human Rights)	~35/100
Ethiopia	National AI strategy finalized 2025 (Diplo)	**Not signed** (Centre for Human Rights)	Ratified Jul 2010 (OHCHR)	Proclamation 568/2008 (Centre for Human Rights)	~30/100
Rwanda	National AI Policy (4 pillars) (ciphersense.ai)	Ratified Jan 2022 (Centre for Human Rights)	Ratified Dec 2008 (OHCHR)	Law 01/2007 (Centre for Human Rights)	~35/100
Mauritius	Blueprint 2025–2029 (ciphersense.ai)	**Not signed** (Centre for Human Rights)	Ratified Jan 2010 (OHCHR)	Bill 2023 (Centre for Human Rights)	~40/100
South Africa	National AI Policy Framework (ciphersense.ai)	Ratified Feb 2023 (Centre for Human Rights)	Ratified Nov 2007 (OHCHR)	Disability Rights Charter (Centre for Human Rights)	~50/100

Table 1: Sub-Saharan African countries with active AI strategies and their disability rights frameworks. GDPPI scores indicate disability inclusion in AI governance. The "High Risk" threshold is 50/100.

The pattern is consistent across all eight countries: **CRPD ratification is nearly universal**, but **ADP ratification is sparse**, **domestic disability legislation is often outdated or unimplemented**, and **none of the national AI strategies include disability accessibility provisions**. Nigeria, despite its relatively strong GDPPI score of 45/100, has a National Commission for Persons with Disabilities that "**often competes with disability-focused CSOs for implementation roles**" and whose engagement with disability organizations is "**weak, limited mostly to occasional event invitations without substantial involvement in programme design or evaluation**" (globaldisabilityfund.org) . Egypt, the continent's AI leader with a "perfect policy score" on AI readiness indices, has **not even signed the African Disability Protocol** and has no ICT accessibility legislation covering government AI systems (ciphersense.ai) .

The launch of the **Hub for AI and Disability Inclusion (HAIDI)** and the **African Disability Data Network (ADDN)** at the Global Data Festival 2026 marks a recognition of this gap, but these initiatives are nascent and funded by external donors (IDRC, FCDO Services, AI4D) rather than by the African governments whose AI strategies created the problem (techreviewafrica.com) . The Assistive Technologies for Disability Trust, which launched HAIDI, described the situation starkly: "**the absence**

of comprehensive disability-related data continues to limit the effectiveness of AI systems used in key sectors such as education, healthcare, employment, and public service delivery across the continent" (techreviewafrica.com) .

6.2 The Middle East and North Africa: High Investment, Low Enforcement

The MENA region combines the highest per-capita government AI investment with significant disability law implementation gaps. Saudi Arabia, the UAE, Qatar, and Egypt have all launched major national AI programs, but the CRPD Committee's concluding observations reveal a consistent pattern: accessibility initiatives exist but are "**not systematic**," "**associated with charitable initiatives**," and "**limited in scope**" (International Disability Alliance) .

Saudi Arabia's GDPPI score of **88/100** — the highest among non-OECD nations — reflects strong AI strategy inclusion and recent disability law reforms. The Authority for the Care of Persons with Disabilities was established, and Vision 2030 includes equity and lifelong learning goals (policyjssr.com) . But the qualitative study of AI healthcare in Saudi Arabia found that "**telemedicine platforms introduced during COVID-19 lack features such as sign language interpretation or text-to-speech capabilities**" and that "**the Sehhaty application... presents navigation challenges for users with visual impairments**" (MDPI) . The country's ALLaM LLM, which will increasingly power government digital services, was not evaluated for accessibility before deployment.

Qatar's **Mada Assistive Technology Center** champions digital inclusion through training and technology provision, and its **JUMLA Sign Language project** uses avatar technology for Arabic Sign Language (resoinsights.com) . But Qatar's AI initiatives, including its partnership with WISE on AI for Accessibility and Inclusion, remain voluntary and project-based rather than embedded in binding procurement requirements or model evaluation standards (Source) . The CRPD Committee has not yet issued concluding observations on Qatar's CRPD compliance, but the country's Law No. 2 of 2004 on the Rights of Persons with Disabilities predates the AI era and does not address algorithmic accessibility.

6.3 Central Asia and the Caucasus: The Emerging Frontier

Central Asia represents the newest frontier of government AI development, with the weakest disability rights infrastructure. Kazakhstan, the regional leader, ratified the CRPD in 2015 and developed KazLLM through its Institute of Smart Systems and Artificial Intelligence (MDPI) . But Uzbekistan, Kyrgyzstan, Tajikistan, and Turkmenistan are only beginning to develop AI strategies, and their disability rights frameworks are even less developed.

Turkmenistan is particularly illustrative. The country ratified the CRPD in 2008 but has **no disability-specific legislation** and ranks among the world's most politically isolated states. In 2026, UNDP began supporting Turkmenistan's Ministry of Communications to develop a **National AI Strategy**, with workshops on "strengthening institutional readiness, advancing digital infrastructure, and identifying priority policy directions" (United Nations Development Programme) . The draft strategy references "inclusive growth and sustainable development" but there is no evidence that disability accessibility has been included in the policy framework. Given Turkmenistan's absence of domestic disability legislation and its isolation from international civil society, the risk that its forthcoming AI systems will exclude disabled citizens is extremely high.

6.4 Southeast Asia: Rapid AI Adoption Without Accessibility Anchors

Southeast Asia's AI governance landscape is characterized by rapid strategy development, ethics-oriented frameworks, and minimal disability-specific provisions. Thailand's National AI Strategy includes genera-

tive AI guidelines mandating risk assessments and transparency, and a royal decree under review adopts a **risk-based model inspired by the EU AI Act** (The National Bureau of Asian Research (NBR)) . But the CRPD Committee found "**the lack of comprehensive and effective measures**" including "**the lack of information and communication technologies and systems**" (International Disability Alliance) . Indonesia's National AI Strategy 2020–2045, embedded in the Making Indonesia 4.0 agenda, has no disability accessibility provisions, and the CRPD Committee noted "**the lack of comprehensive national accessibility legislation and strategy across all levels of government**" including "**information and communications technology**" (International Disability Alliance) . The Philippines' Roadmap 2.0 aims to make the country a regional AI Center of Excellence, but existing accessibility laws "**do not include principles of accessibility through universal design**" (International Disability Alliance) .

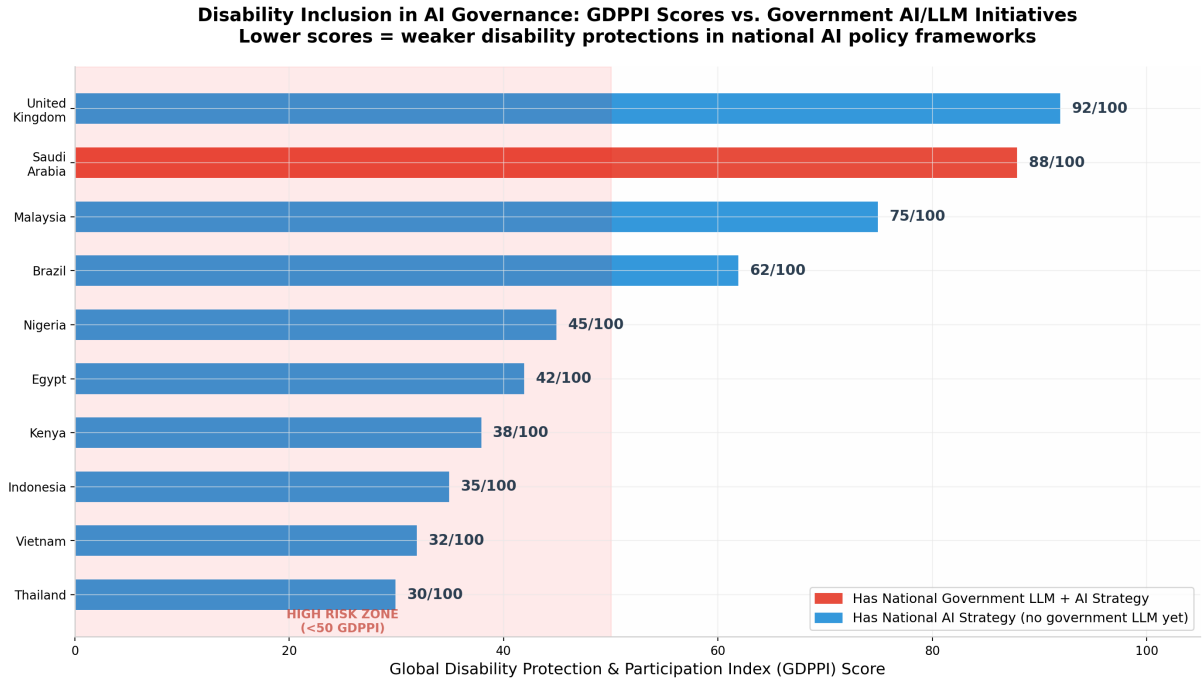

Figure 3: GDPPI Scores Chart

Figure 3: Global Disability Protection & Participation Index (GDPPI) scores for selected countries with government AI/LLM initiatives. The red zone (<50 GDPPI) marks jurisdictions where disability inclusion in AI governance is critically weak. Saudi Arabia, despite having a national LLM (ALLaM), scores 88/100 due to Vision 2030 disability reforms — but its AI model was still not evaluated for accessibility.

---

7. The Human Cost: What Algorithmic Exclusion Looks Like in Practice

7.1 The Screen Reader Failure

When a government deploys an LLM-powered chatbot on its public services portal, blind users interact with it through screen readers — software that converts visual interface elements into speech or Braille output. If the chatbot's interface uses semantic HTML incorrectly, fails to update ARIA live regions when new messages arrive, or includes buttons without accessible labels, the screen reader cannot convey the conversation to the user. A 2026 analysis of AI chatbot accessibility identified common failures including "**focus jumps to the top of the page after a message is submitted**," "**buttons and icons with**

no accessible labels," "**new messages not announced automatically**," and "**excessive emojis that each read out as their full Unicode description**" (Enabled.in) .

These are not hypothetical concerns. New Zealand's 2026 government accessibility audit found that disabled citizens regularly encounter "**inaccessible security processes**" including "**complex password requirements and frequent log-outs that disproportionately affect those with learning disabilities, motor impairments, or who use assistive technology**" and "**CAPTCHA systems — both visual and auditory**" that are "**particularly challenging for blind users**" (dns.govt.nz) . One blind participant reported: "**When you have to identify you're not a robot with pictures — especially when you get to the end of a long form... I've cried in the past. Sometimes they have an audio option, but that's hard**" (dns.govt.nz) . When these same citizens interact with a government LLM that generates content incompatible with their assistive technologies, the exclusion is compounded.

7.2 The Speech Recognition Gap

For citizens with speech impairments, LLM-powered voice interfaces present a double barrier. The model must both understand the user's speech and generate responses that the user can consume through their preferred modality. The Saudi healthcare study documented this precisely: a participant with a speech impairment reported that "**the AI speech recognition system struggles with my speech patterns. Despite being marketed specifically for speech impairments, it requires constant retraining to understand my voice**" (MDPI) . The Penn State BITS study found that "**tools that predict aptitude and ability based on speech patterns do not account for people with speech and communication disabilities**" and that "**some chatbots may be inaccessible based on visual or other processing or communication disabilities**" (National Disability Institute) .

Government LLMs trained on standard speech datasets — which predominantly feature neurotypical, able-bodied speakers — will reproduce these failures at scale. The AccessEval benchmark found that **speech impairments were among the most affected disability types** in LLM bias evaluation, with models producing "**more factually incorrect and denial responses**" for speech-related queries (ACL Anthology) .

7.3 The Cognitive Accessibility Deficit

Cognitive disabilities — including intellectual disabilities, autism, ADHD, and learning disabilities — present distinct accessibility challenges that government LLMs are particularly ill-equipped to address. LLMs typically generate long, complex, jargon-heavy responses that overwhelm users with cognitive disabilities. A 2026 analysis found that "**long, complex, jargon-heavy responses can overwhelm users, especially those with cognitive disabilities, ADHD, learning disabilities, or low literacy levels**" and that "**session timeouts**" create additional stress when "**a user needs more time to read or respond**" (Enabled.in) .

The AccessEval study found that LLMs frequently produce "**misleading, less informative, overly cautious, or overly general responses**" for disability-related queries, "**often not providing disability-specific information**" and "**incorrectly suggesting inapplicable technologies that do not directly assist with a given disability**" (ACL Anthology) . For a user with an intellectual disability seeking government services through an LLM interface, these failures are not inconveniences — they are barriers to exercising fundamental rights.

---

8. The Path Forward: Binding Obligations for Accessible Government AI

8.1 What International Law Already Requires

The existing international legal framework is not silent on this issue — it is simply **unenforced**. The CRPD's Article 9 obligations are clear and binding on 185 states. General Comment No. 2 explicitly covers "**all ICT products and ICT based applications and services**" (G3ict) . The Marrakesh Treaty creates obligations for accessible format reproduction that apply to AI-generated content. The African Disability Protocol, once in force, will add a continental layer of obligation for AU member states. The UN Special Rapporteur on the rights of persons with disabilities has called for "**persons with disabilities and their representative organizations**" to be "**involved in the development, procurement and deployment of AI systems and in the policy-making process**" as a requirement of the CRPD (European Disability Forum) .

The NYC Bar Association's 2025 recommendations for AI and disability rights provide a concrete policy framework that governments could adopt immediately: **mandate full accessibility in AI development**; **require the inclusion and weighting of disability data** in AI training sets; **establish standardized, sector-specific accessibility benchmarks** for generative and assistive AI evaluations; **support hybrid AI architectures** combining deterministic accessibility rules with machine learning; and **guarantee non-AI alternatives and human-centric options** for all AI-driven services (New York City Bar Association) .

8.2 The Minimum Viable Standards for Government LLMs

Any government deploying an LLM as a public service should be required to meet the following minimum standards, derived from the CRPD, WCAG 2.1, and emerging best practices in AI accessibility:

Requirement	Legal Basis	Implementation
Training data must include disability-representative content	CRPD Art. 4(1)(f) —R&D of universally designed goods (OHCHR) ; NYC Bar Rec. 3 (New York City Bar Association)	Minimum 5% of training corpus should be disability-related content; annotate for disability type
Model evaluation must include disability-disaggregated metrics	CRPD Art. 9(2)(a) —monitor accessibility implementation (un.org) ; AccessEval framework (ACL Anthology)	Test on disability-specific query sets; report performance by disability type
Interface must comply with WCAG 2.1 AA	CRPD Art. 9(1) —ICT accessibility (un.org) ; WEF/ ITU Model Policy (World Economic Forum)	Screen reader compatibility, keyboard navigation, color contrast, alt-text
Multi-modal input/output must be supported	CRPD Art. 9(2)(e) —live assistance and intermediaries (un.org)	Text, voice, and visual alternatives for all interactions
Procurement must include accessibility criteria	CRPD Art. 9(2)(a); Committee recommendations to UAE, Poland, Hungary (International Disability Alliance)	Mandatory accessibility clauses in all AI procurement contracts

Table 2 – continued

Requirement	Legal Basis	Implementation
Human alternative must always be available	NYC Bar Rec. 7 (New York City Bar Association)	Non-AI option for all AI-driven services, with equivalent convenience and timeliness
DPO participation in design and evaluation	CRPD Art. 4(3) —consultation with PWDs (OHCHR) ; UN SR Recommendation (European Disability Forum)	Organizations of persons with disabilities involved at all stages
Model card must include accessibility statement	Emerging best practice (Human in the Loop)	Published documentation of accessibility features, limitations, and known failures

Table 2: Minimum viable accessibility standards for government-deployed LLMs, with legal basis and implementation guidance.

8.3 The Institutional Reforms Required

Closing the sovereign AI–disability rights gap requires institutional reforms at three levels. **First**, national AI strategies must be amended to include dedicated disability accessibility pillars, with budget allocations, procurement requirements, and evaluation metrics. The AU's Continental AI Strategy should be revised to include a disability inclusion mandate for all member state AI programs. **Second**, disability rights agencies must be given formal authority to review and approve government AI procurement decisions, with the power to block deployments that fail accessibility standards. Nigeria's NCPWD should have veto authority over NCAIR procurement that lacks accessibility criteria; Saudi Arabia's Authority for the Care of Persons with Disabilities should review ALLaM deployments. **Third**, international donors and development banks — including the World Bank, which actively supports digital public infrastructure and AI readiness programs — must make accessibility compliance a condition of funding, as the World Bank has already begun to do for general digital inclusion (The World Bank) .

---

9. Conclusion: Sovereignty Cannot Mean Exclusion

The global movement toward sovereign AI is, in many respects, a positive development. It reduces dependence on foreign technology, preserves linguistic and cultural diversity in AI systems, and creates opportunities for locally relevant innovation. But sovereignty without accessibility is not liberation — it is a transfer of exclusion from foreign corporations to domestic governments. When a state builds an LLM and deploys it as public infrastructure without ensuring that disabled citizens can use it, the state is not exercising digital sovereignty; it is **reproducing discrimination at scale**.

The evidence assembled in this exposé demonstrates that this is not a marginal problem affecting a few jurisdictions. It is a **systemic, global failure** that cuts across every region where government AI initiatives are active. Saudi Arabia's ALLaM, trained on 500 billion tokens without disability data, evaluated by 400 experts without disability expertise, and deployed on platforms without accessibility testing. Nigeria's AI Scaling Hub, launched with a GDPPI score of 45/100 and no disability procurement requirements. Egypt's Second National AI Strategy, with six pillars and zero disability provisions. Turkmenistan's forthcoming National AI Strategy, being developed in a country with no disability legislation

and no civil society input. The African Union's Continental AI Strategy, calling for "African AI" while 75% of its member states have not ratified the African Disability Protocol.

The international legal framework to prevent this exclusion already exists. The CRPD's Article 9 obligations are binding. The Marrakesh Treaty's accessible format requirements are ratified by 93 countries. The CRPD Committee's concluding observations provide a detailed roadmap of what compliant practice looks like. What is missing is not law but **enforcement, institutional coordination, and political will**. The half-dozen companies whose frontier systems underpin most of the AI economy — OpenAI, Anthropic, Google DeepMind, Meta, Microsoft, and Nvidia — "**have the power**" to change how models are trained, evaluated, and deployed, but "**they do not have the incentive**" because "**the revenue upside... is real but small compared to the upside of the next efficiency frontier**" (Human in the Loop) . Governments, as both regulators and the largest procurers of AI systems, are the only actors with both the power and the obligation to close this gap.

The question is whether they will do so before the infrastructure of algorithmic exclusion becomes too entrenched to dismantle. Every government LLM deployed without accessibility evaluation, every national AI strategy written without disability consultation, every procurement contract signed without accessibility criteria — each of these is a brick in a wall that will take years to tear down. The 1.3 billion people with disabilities worldwide do not have years to wait. The CRPD, now in force for nearly two decades, was meant to ensure that they would never have to.

This exposé is based on documentary research of CRPD Committee concluding observations, national AI strategy documents, LLM model cards and technical papers, academic studies on disability bias in AI, international treaty ratification databases, and development agency reports. All factual claims are sourced to publicly available documents.

Generated by Kimi.ai